
The forces behind the changing Internet: IXPs, content delivery, and virtualization

Prof. Steve Uhlig

Head of Networks research group

Queen Mary, University of London

steve@eeecs.qmul.ac.uk

<http://www.eecs.qmul.ac.uk/~steve/>

Guest professor, Institute of Computing Technology, Chinese
Academy of Sciences

Joint work with:

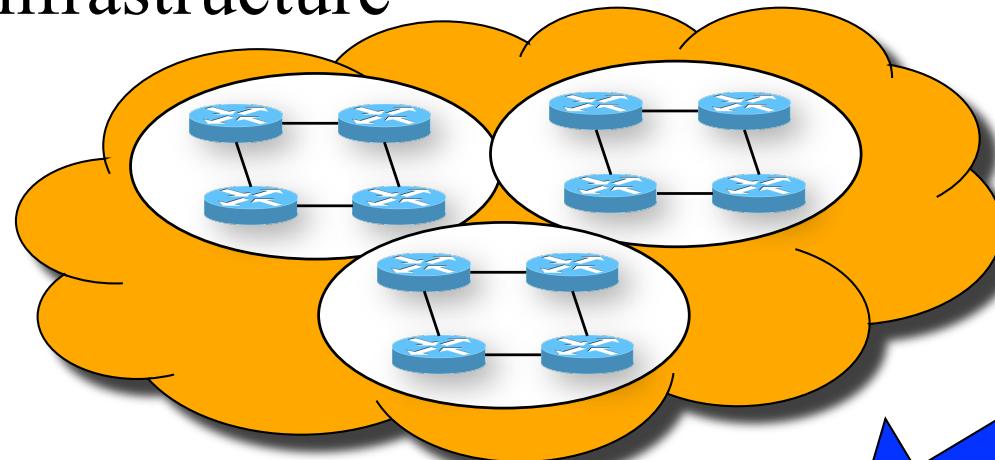
B. Ager (ETH), N. Chatzis, A. Feldmann, B. Frank, I. Poese, N. Sarrar, G. Smaragdakis
(T-labs), Bruce Maggs (Duke/Akamai), Walter Willinger (AT&T), Gaogang Xie (CAS)

Internet: communication platform

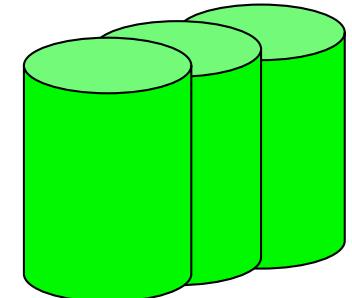
Internet: content sharing and
delivery platform

Content-Processing-Network

Infrastructure

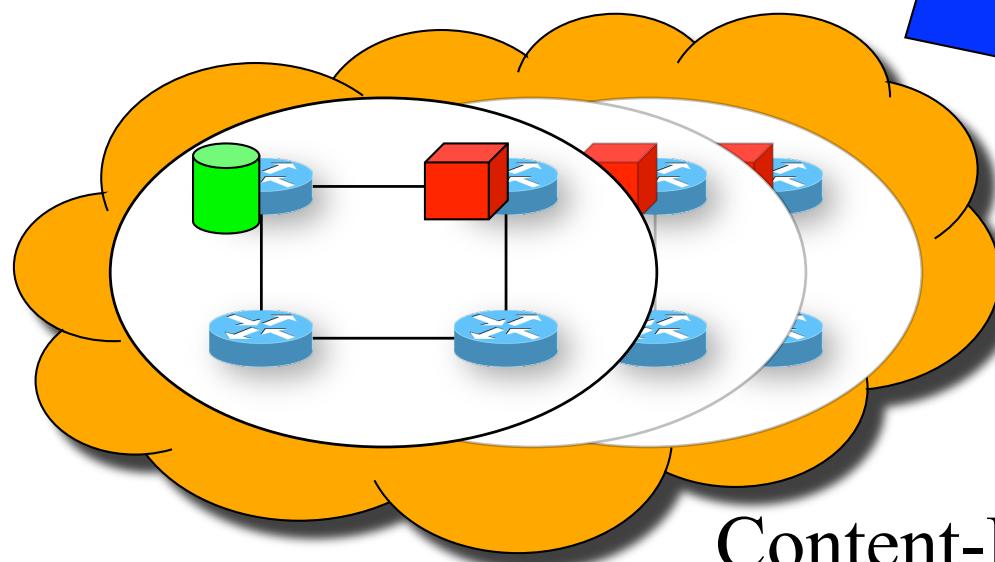
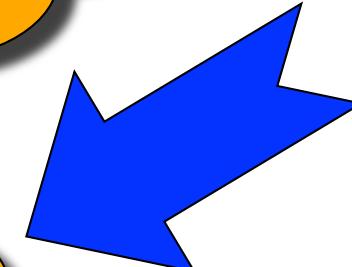
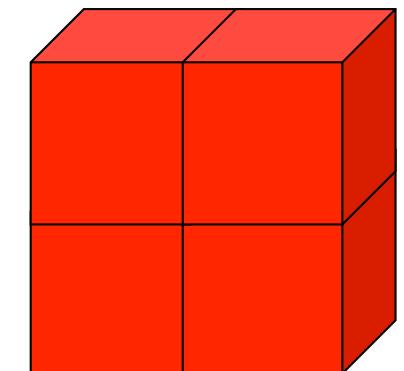


Storage



+

Processing

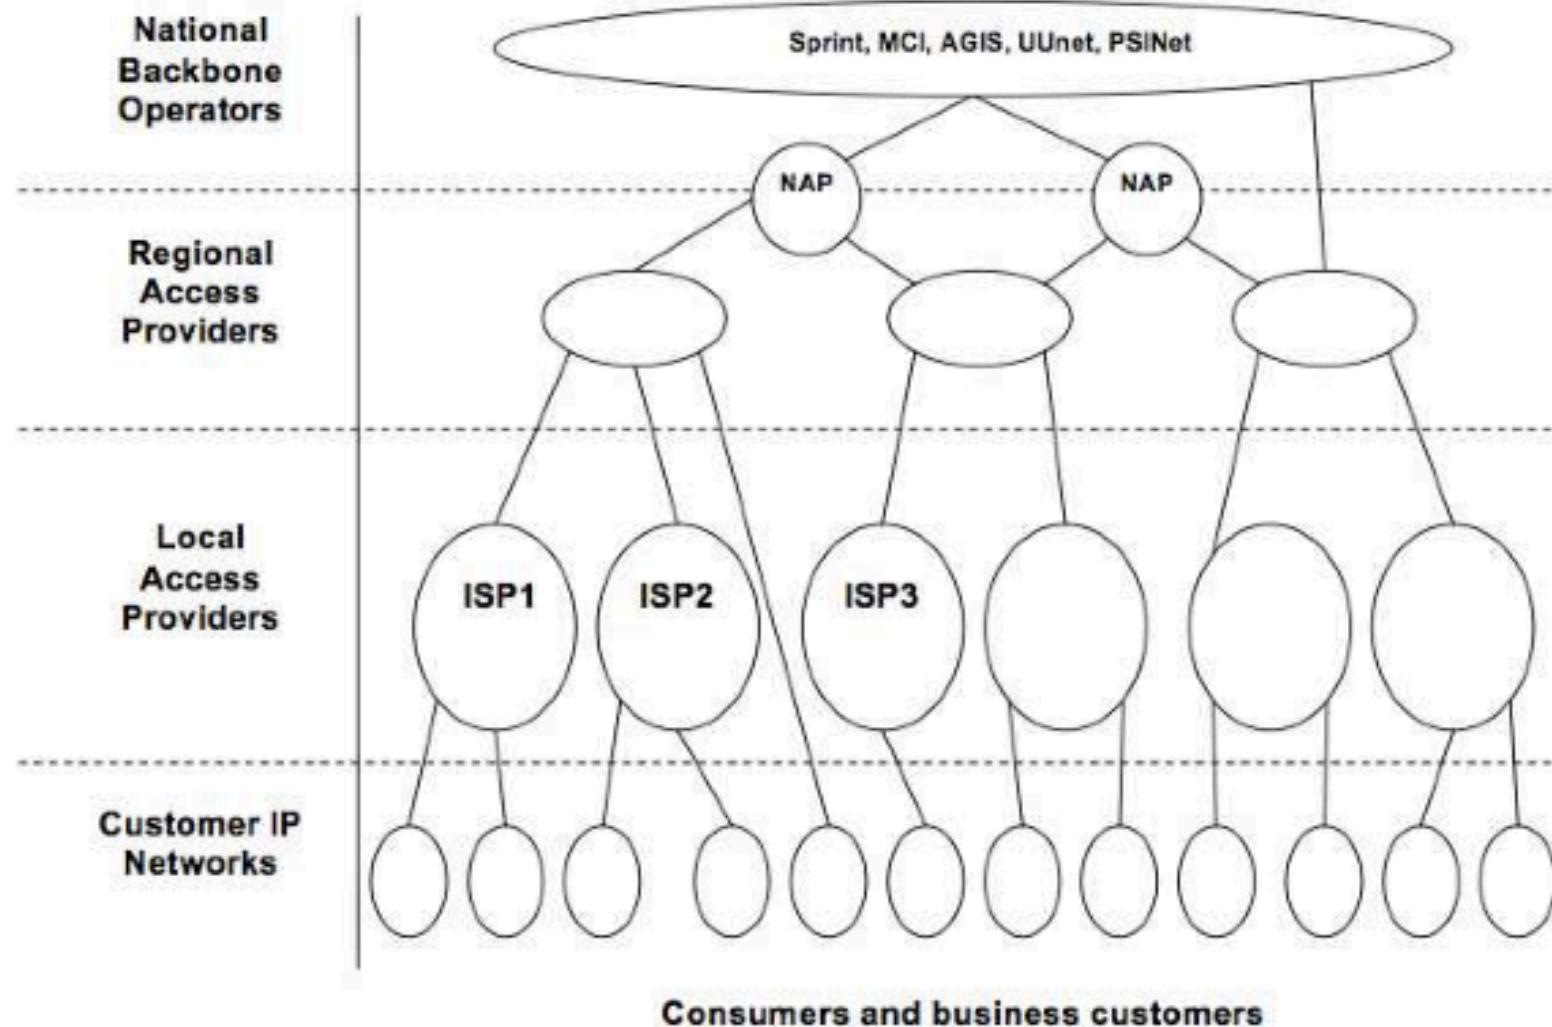


Content-Processing-Network

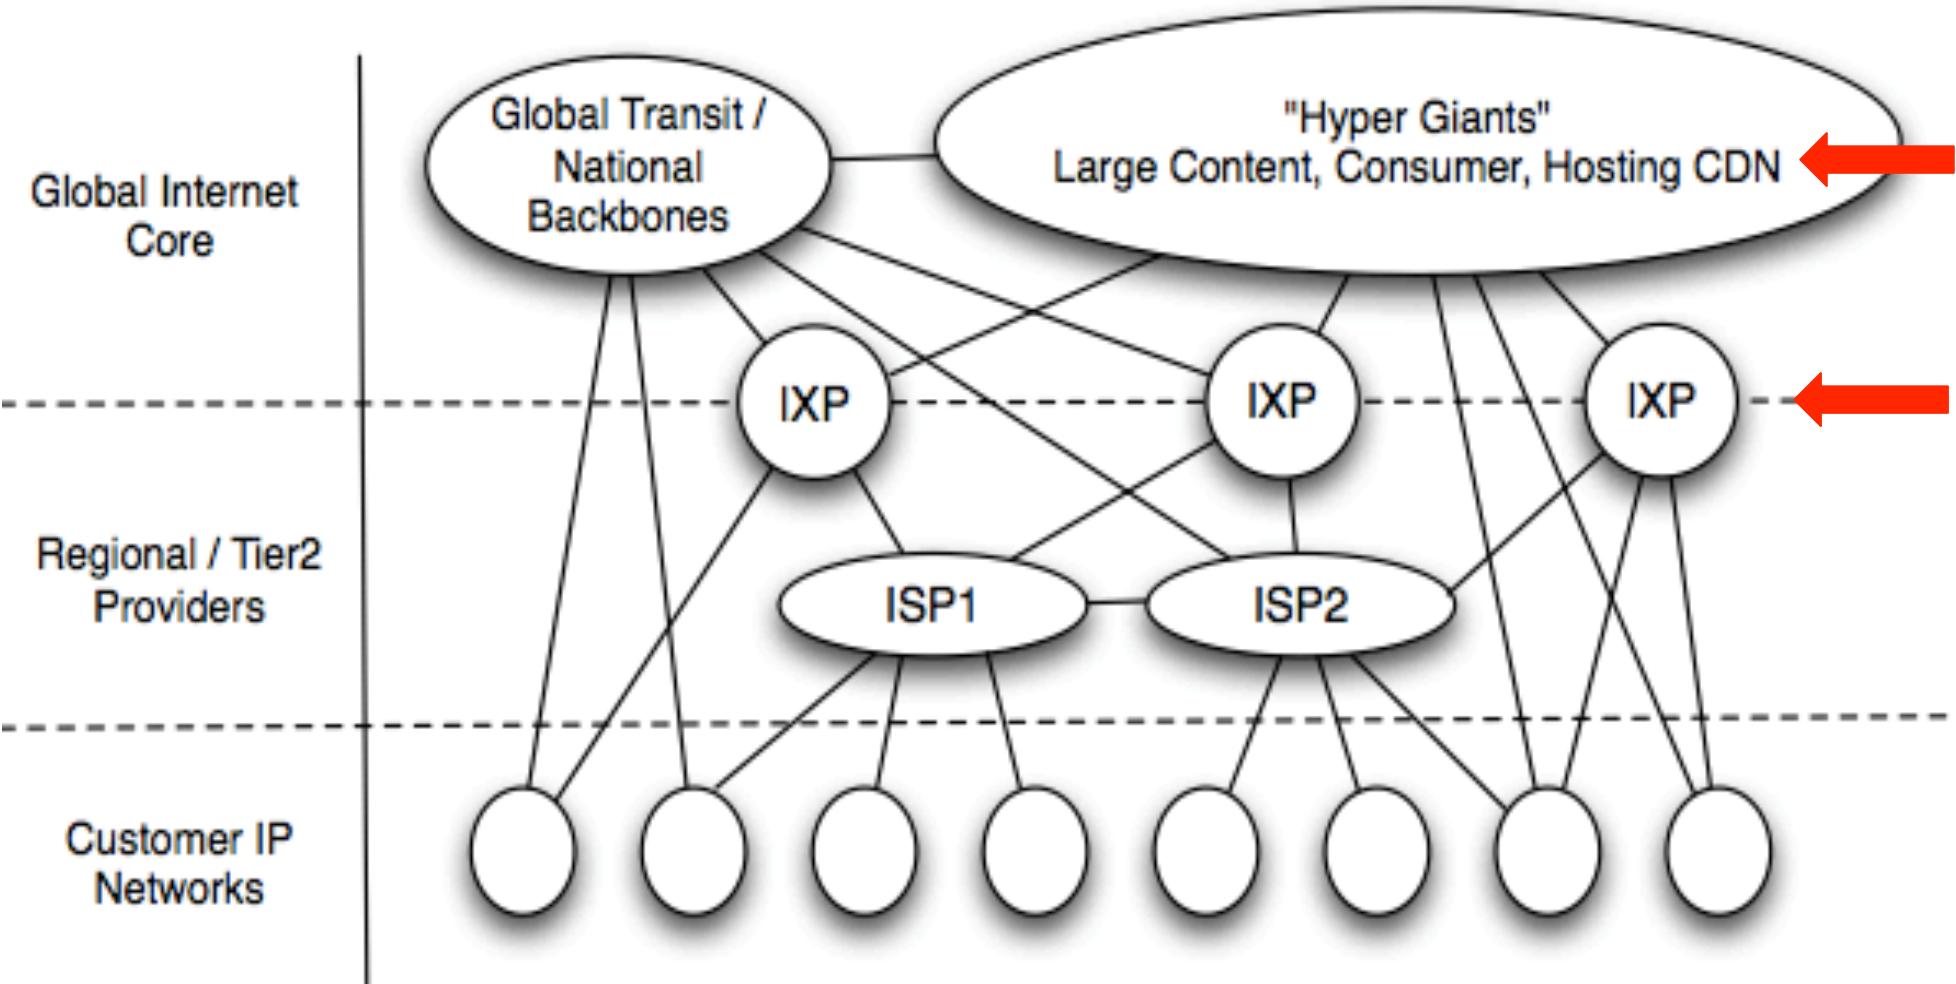
Agenda

- Internet update
- Short history of content delivery
- Content delivery today
- Network virtualization
- Challenges & opportunities

Old mental model



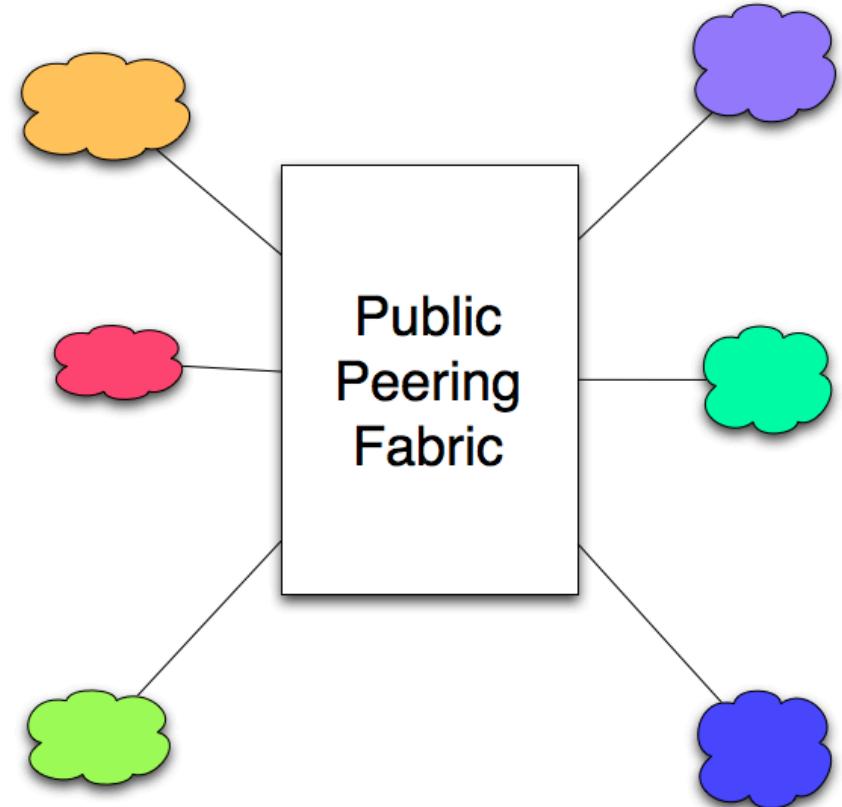
The Internet is flat



C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. *Internet Interdomain Traffic*. Proc. of ACM SIGCOMM, 2010.

Internet Exchange Point

- An Internet exchange point (IXP) is a layer 2 service to facilitate the interconnection between
 - ISPs
 - Hosting or service providers
 - CDNs
- An IXP facilitates peering between players, usually across a public and/or private peering fabric of some type
- Offer public and/or private peerings



B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. *Anatomy of a Large European IXP*. Proc. of ACM SIGCOMM, 2012.

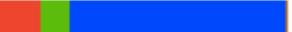
Internet power: from ISP to content

Rank	CAIDA-degree	CAIDA-cone	Renesys	Knodes	Arbor	Potential	Normalized potential
1	Level 3	Level 3	Level 3	Level 3	Level 3	NTT	Chinanet
2	Cogent/PSI	AT&T	Global Crossing	Cogent	Global Crossing	Tinet	Google
3	AT&T	MCI	Sprint	Global Crossing	Google	Global Crossing	ThePlanet
4	MCI	Cogent/PSI	NTT	Sprint	*	Deutsche Telekom	SoftLayer
5	Hurricane	Global Crossing	Savvis	Tinet	*	KDDI	China169 backbone
6	Qwest	Sprint	TeliaSonera	NTT	Comcast	Telia	Level 3
7	Sprint	Qwest	Tinet	AT&T	*	Akamai	Rackspace
8	Global Crossing	Hurricane Electric	Verizon	Swisscom	*	Bandcon	China Telecom
9	tw telecom	tw telecom	AT&T	Hurricane	*	Cable and Wireless	1&1 Internet
10	INIT7	TeliaNet	China Telecom	Telia	*	Qwest	OVH

- CAIDA: BGP-based degree or customer-cone
- Renesys: variant of CAIDA-cone
- Knodes: Fixedorbit.com centrality metric
- Arbor: Interdomain traffic
- Potential: hostname-based
- Normalized potential: weighted hostnames

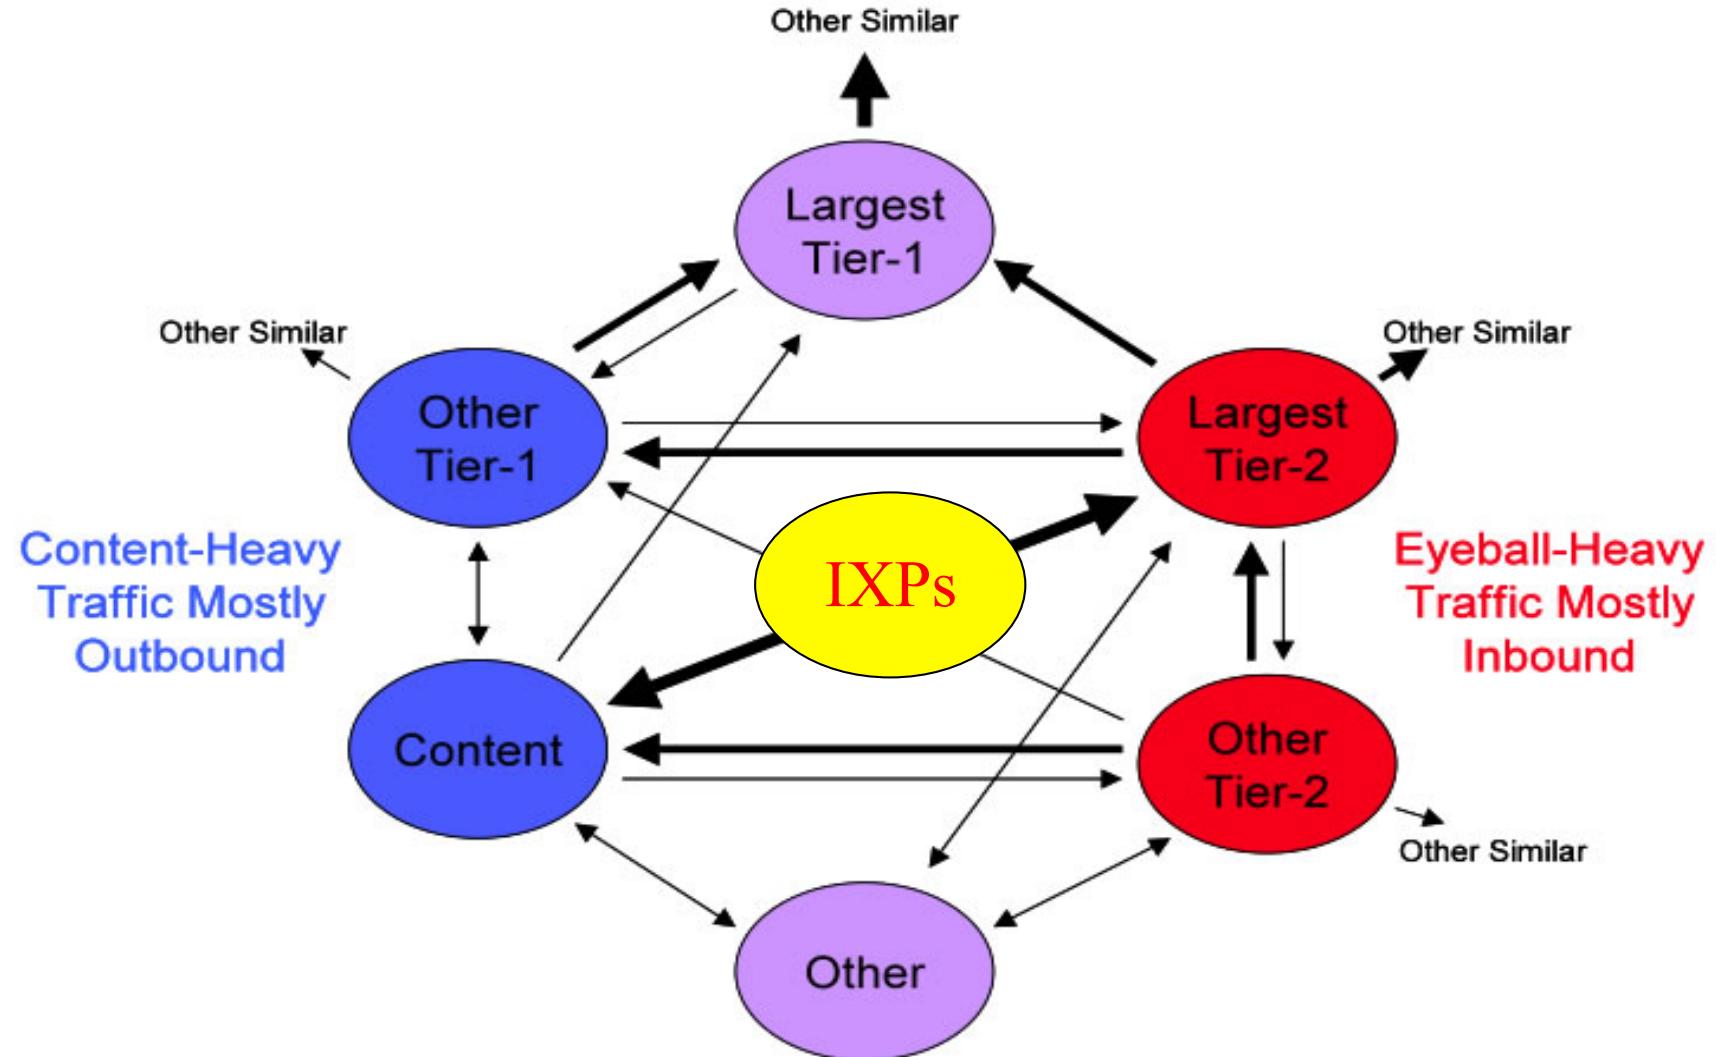
B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig. *Web content cartography*. Proc. of ACM SIGCOMM Internet Measurement Conference, 2011.

Western CDNs

Rank	# hostnames	# ASes	# prefixes	owner	content mix
1	476	79	294	Akamai	
2	161	70	216	Akamai	
3	108	1	45	Google	
4	70	35	137	Akamai	
5	70	1	45	Google	
6	57	6	15	Limelight	
7	57	1	1	ThePlanet	
8	53	1	1	ThePlanet	
9	49	34	123	Akamai	
10	34	1	2	Skyrock	
11	29	6	17	Cotendo	
12	28	4	5	Wordpress	
13	27	6	21	Footprint	
14	26	1	1	Ravand	
15	23	1	1	Xanga	
16	22	1	4	Edgecast	
17	22	1	1	ThePlanet	
18	21	1	1	ivwbox.de	
19	21	1	5	AOL	
20	20	1	1	Leaseweb	

■ only on TOP, ■ both on TOP and EMBEDDED, ■ only on EMBEDDED, ■ TAIL.

Squeezing the ISP

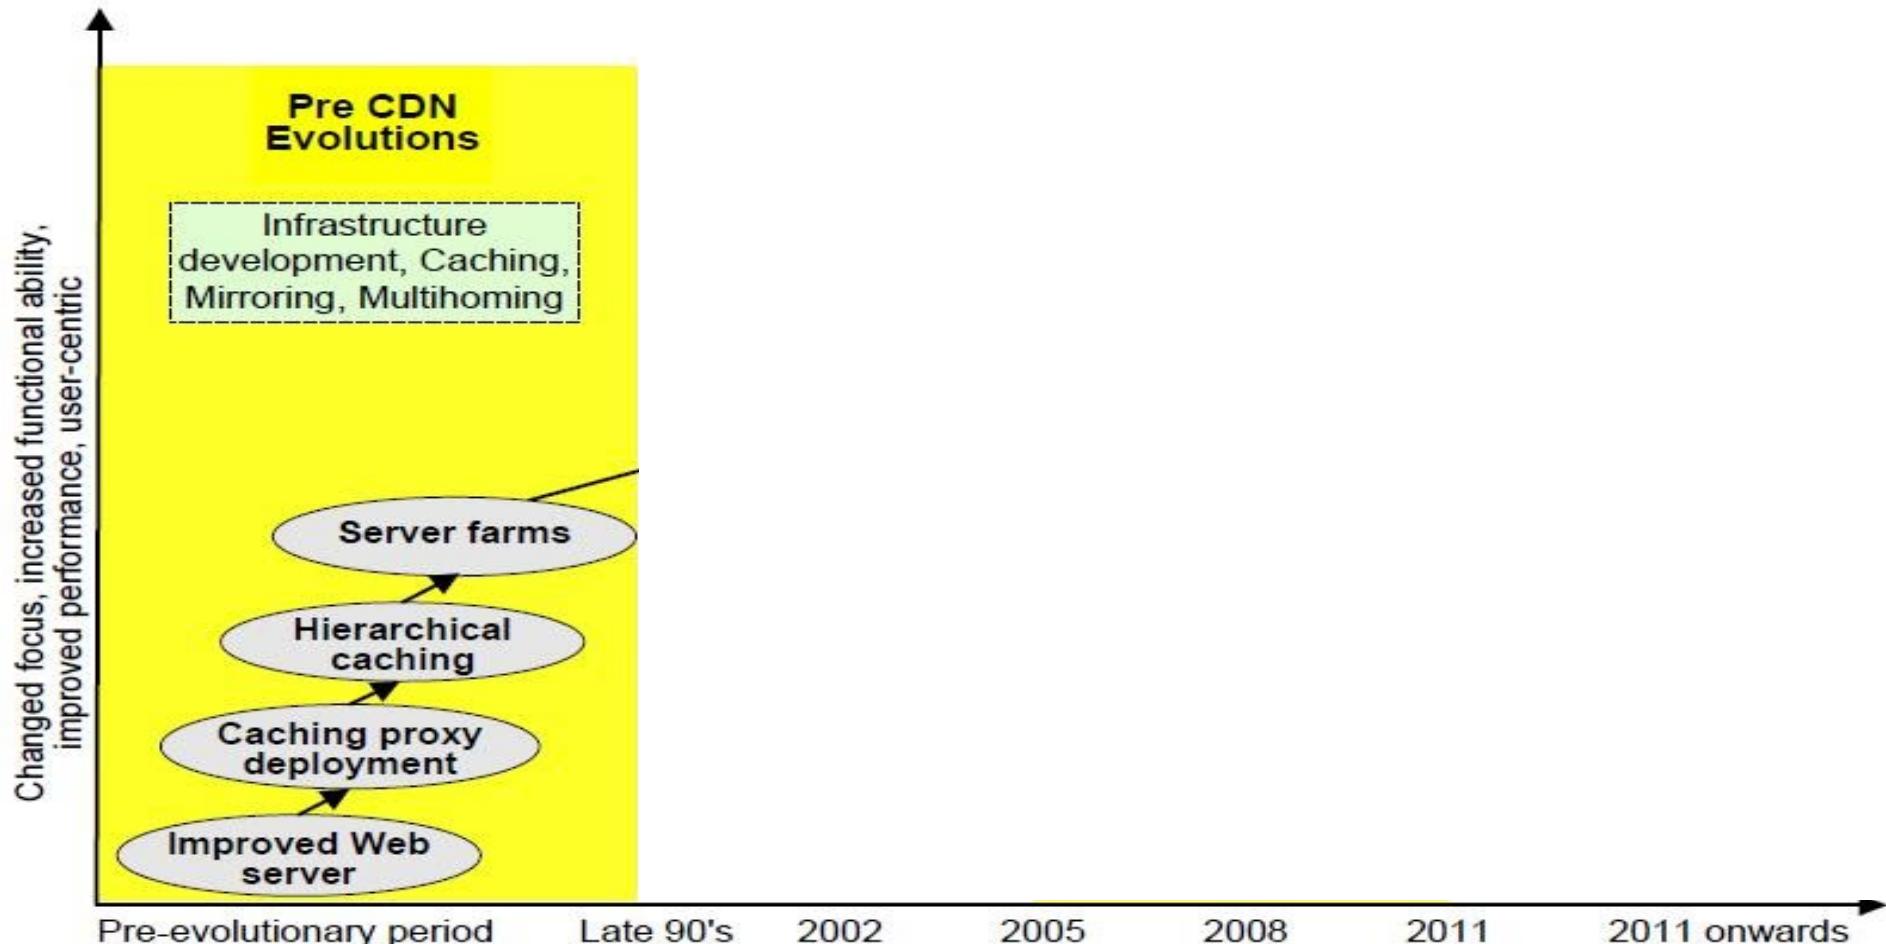


William B. Norton. *The Internet Peering Playbook : Connecting to the Core of the Internet*. DrPeering Press, 2012.

Agenda

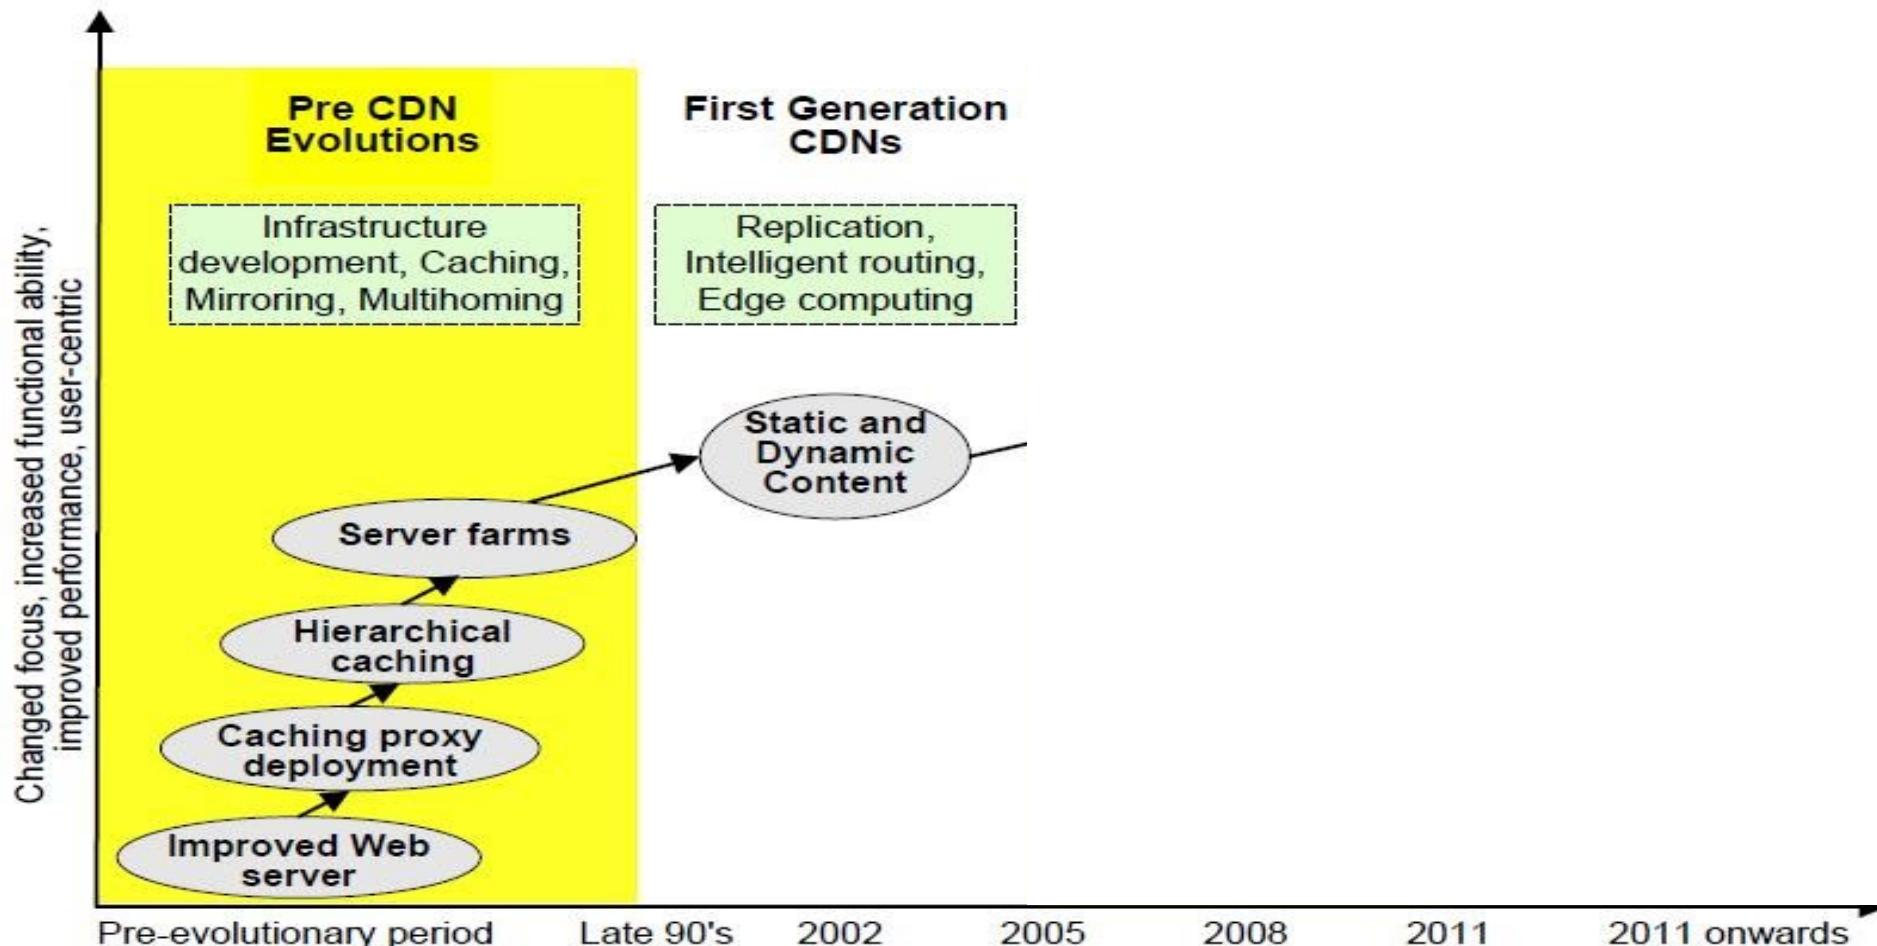
- Internet update
- **Short history of content delivery**
- Content delivery today
- Network virtualization
- Challenges & opportunities

The Early Web



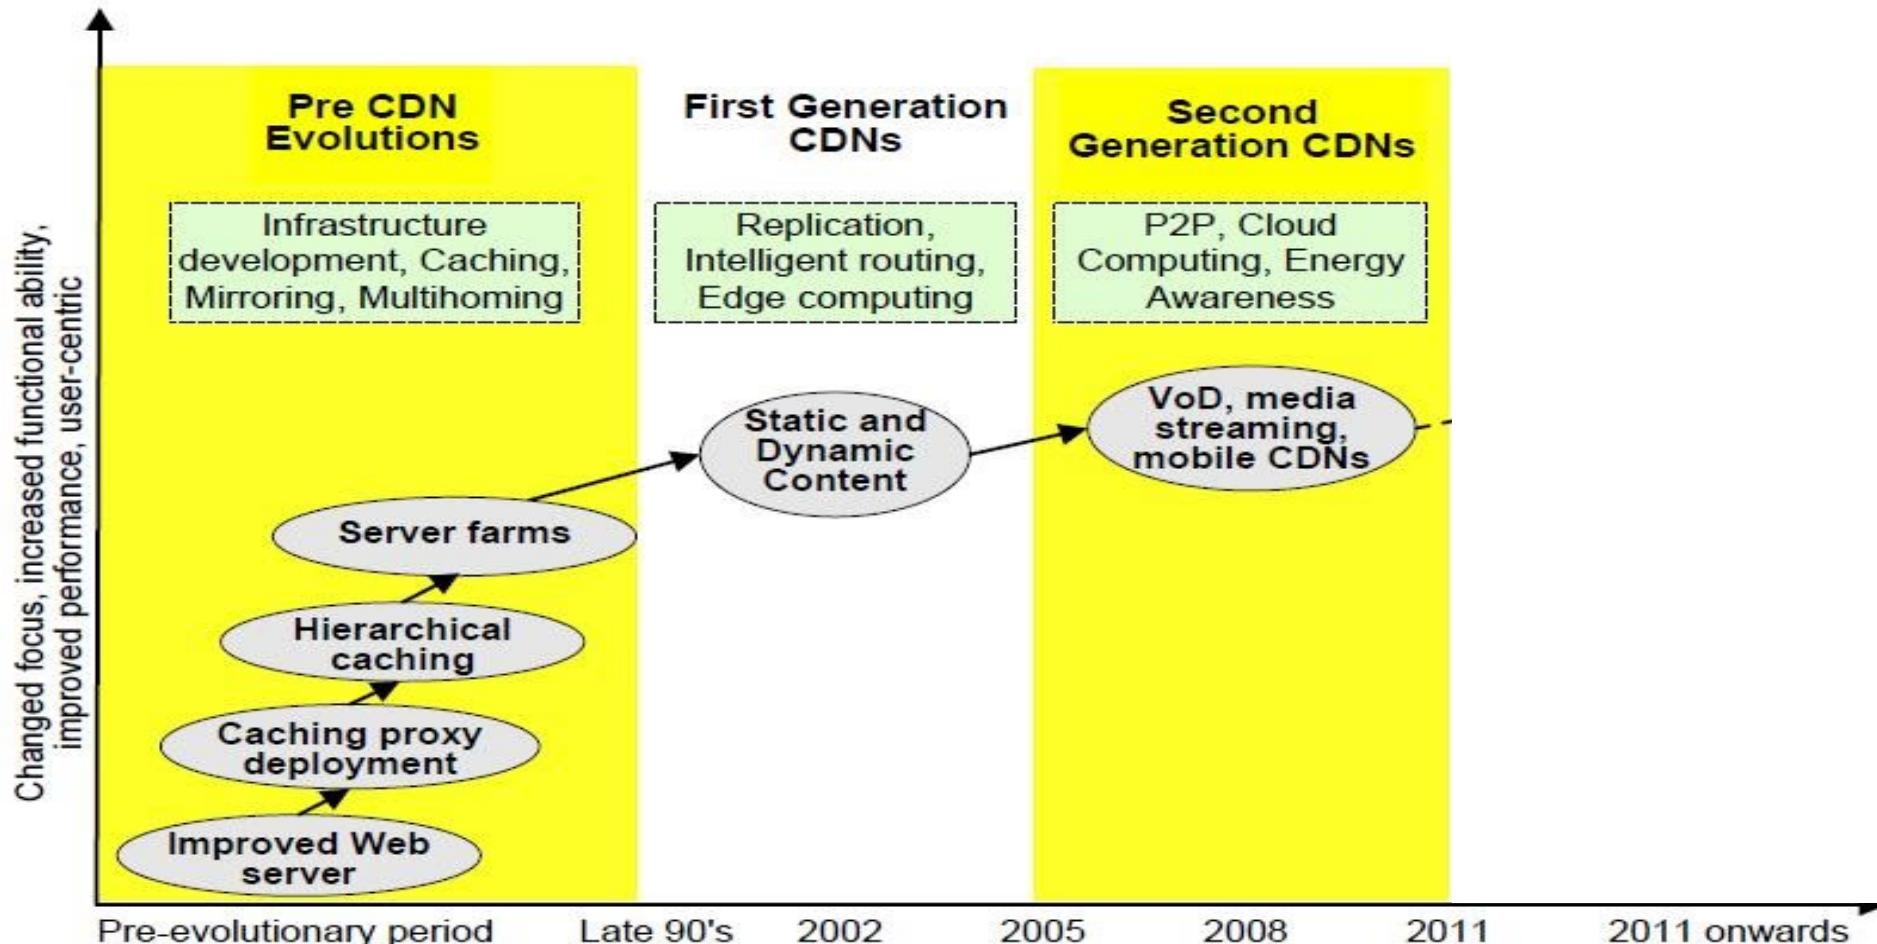
Pathan Mukaddim. Ongoing Trends and Future Directions in Content Delivery Networks (CDNs). Available online from: <http://amkpathan.wordpress.com/article/ongoing-trends-and-future-directions-in-3uxfz2bz8z1w-2/>

CDNs 1.0



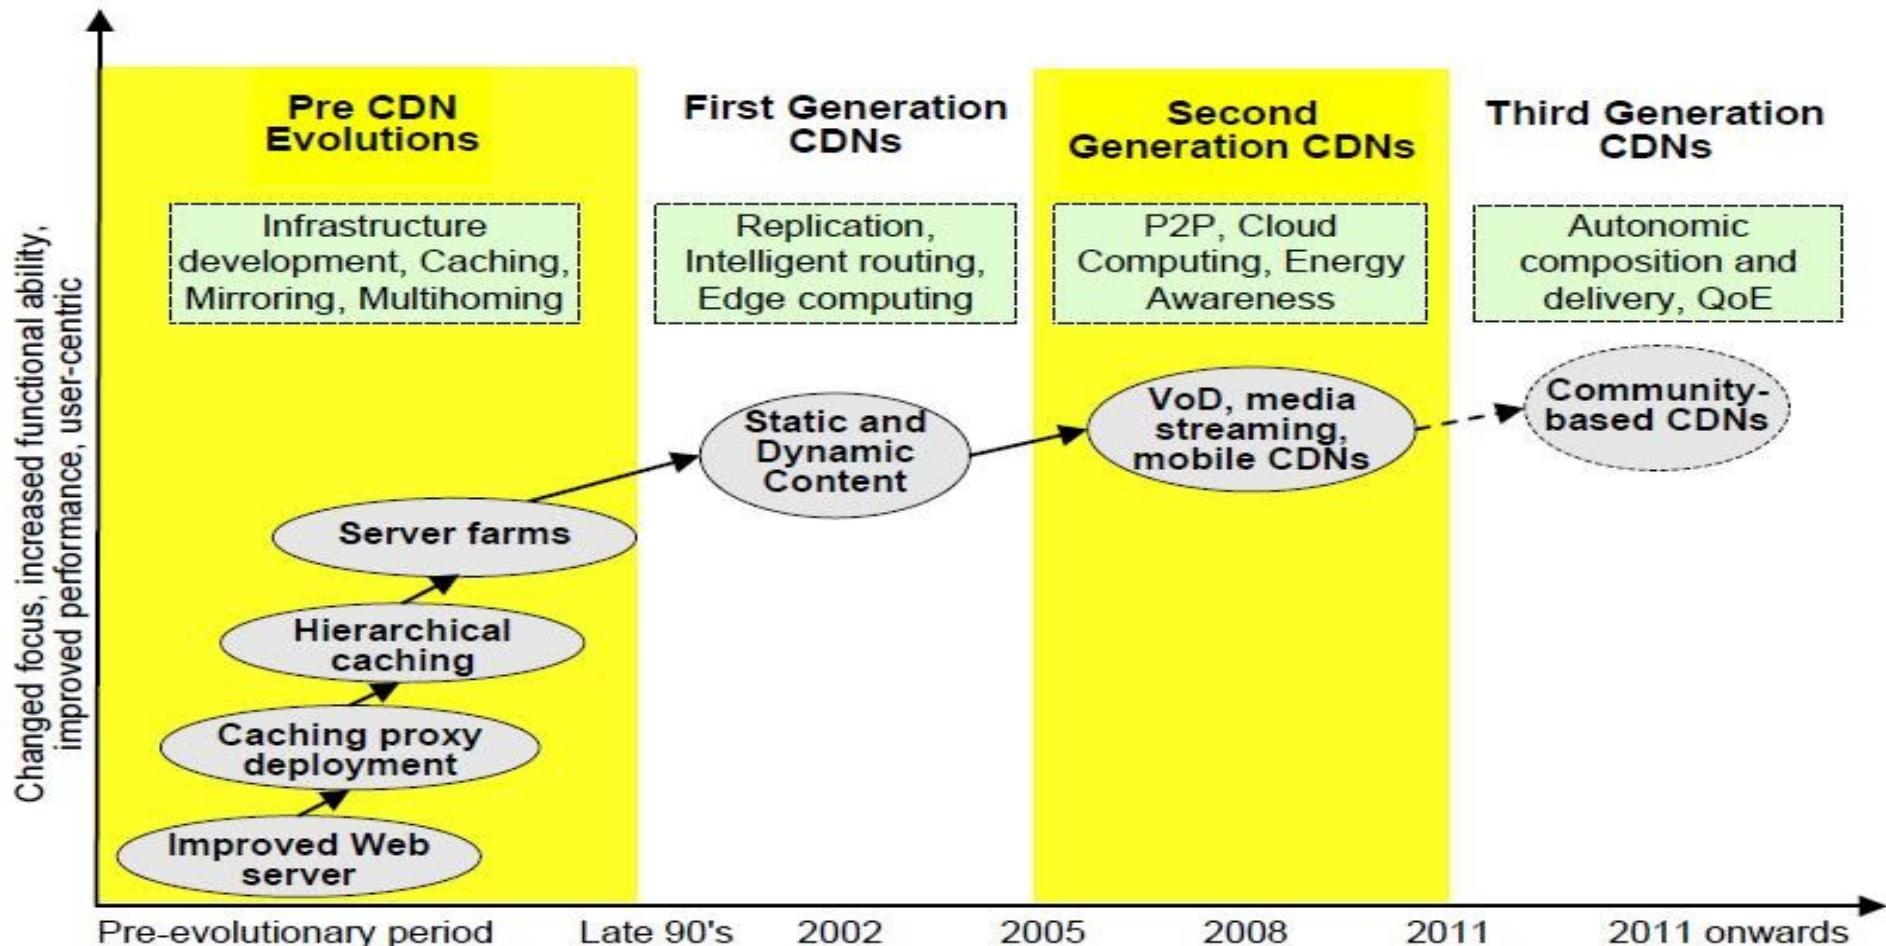
Pathan Mukaddim. Ongoing Trends and Future Directions in Content Delivery Networks (CDNs). Available online from: <http://amkpathan.wordpress.com/article/ongoing-trends-and-future-directions-in-3uxfz2bz8z1w-2/>

CDNs 2.0



Pathan Mukaddim. Ongoing Trends and Future Directions in Content Delivery Networks (CDNs). Available online from: <http://amkpathan.wordpress.com/article/ongoing-trends-and-future-directions-in-3uxfz2bz8z1w-2/>

Autonomic CDNs

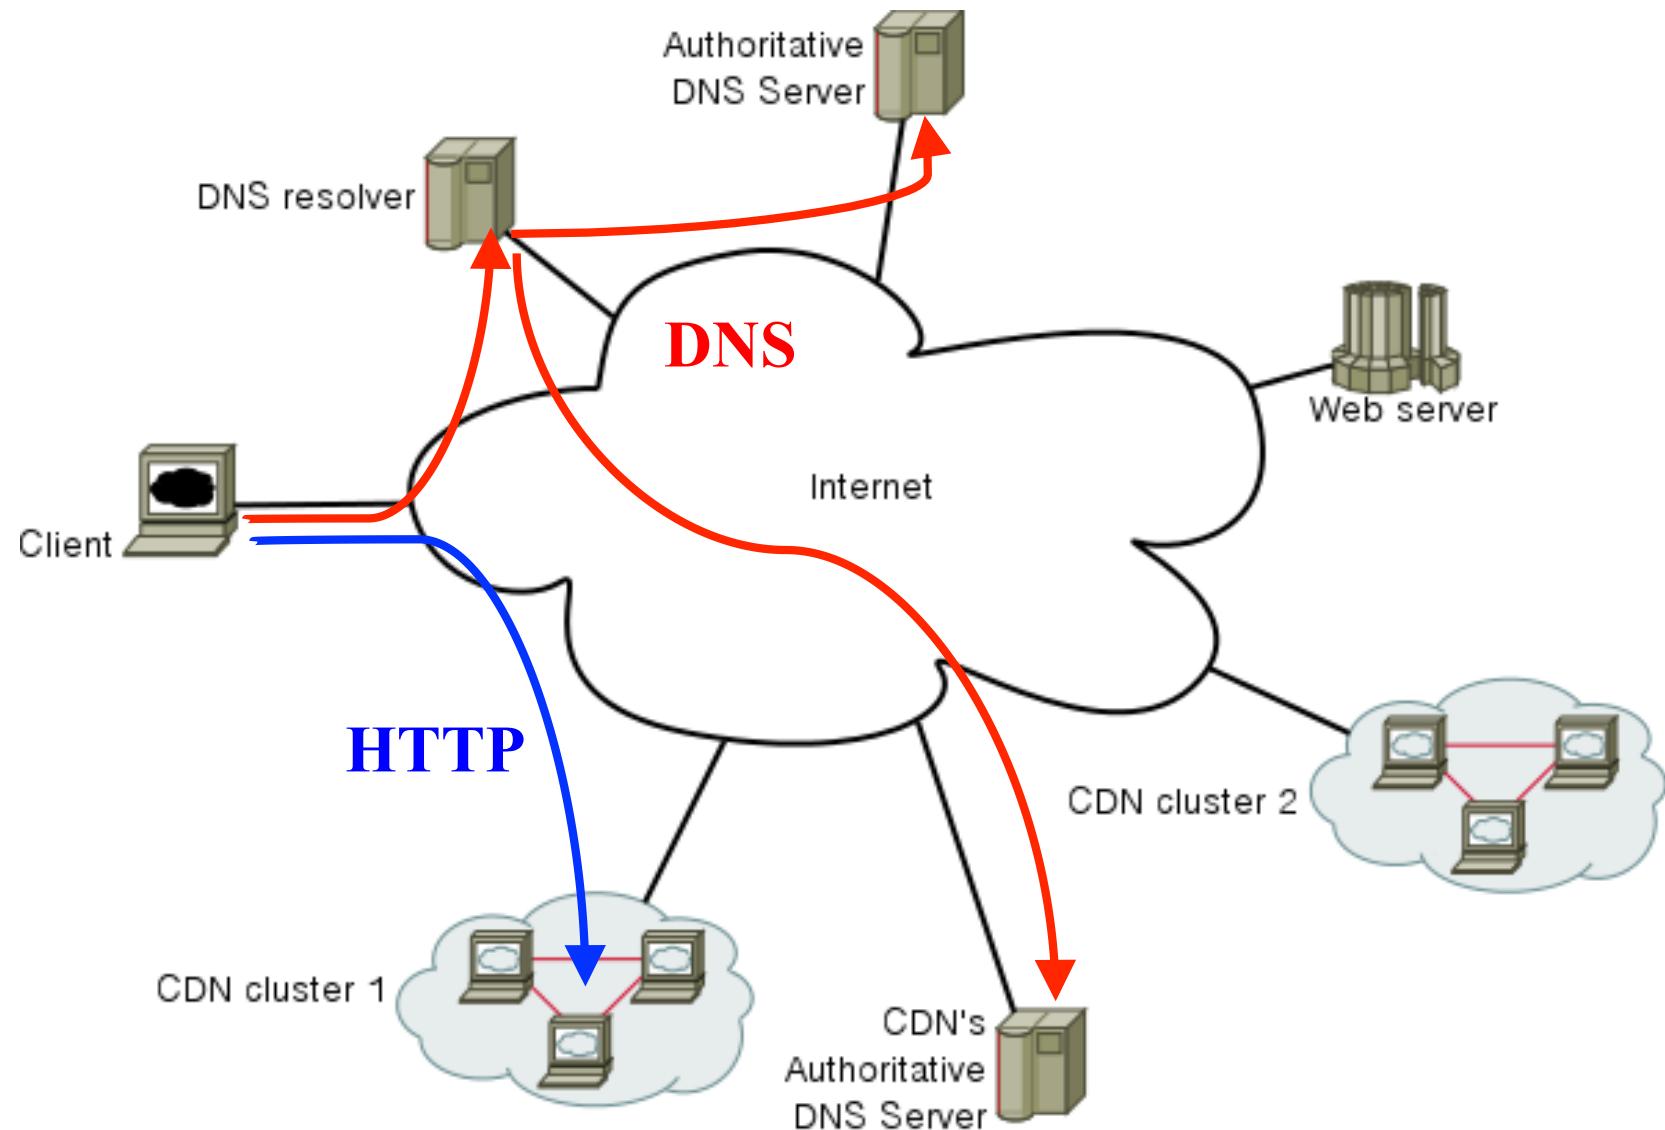


Pathan Mukaddim. Ongoing Trends and Future Directions in Content Delivery Networks (CDNs). Available online from: <http://amkpathan.wordpress.com/article/ongoing-trends-and-future-directions-in-3uxfz2bz8z1w-2/>

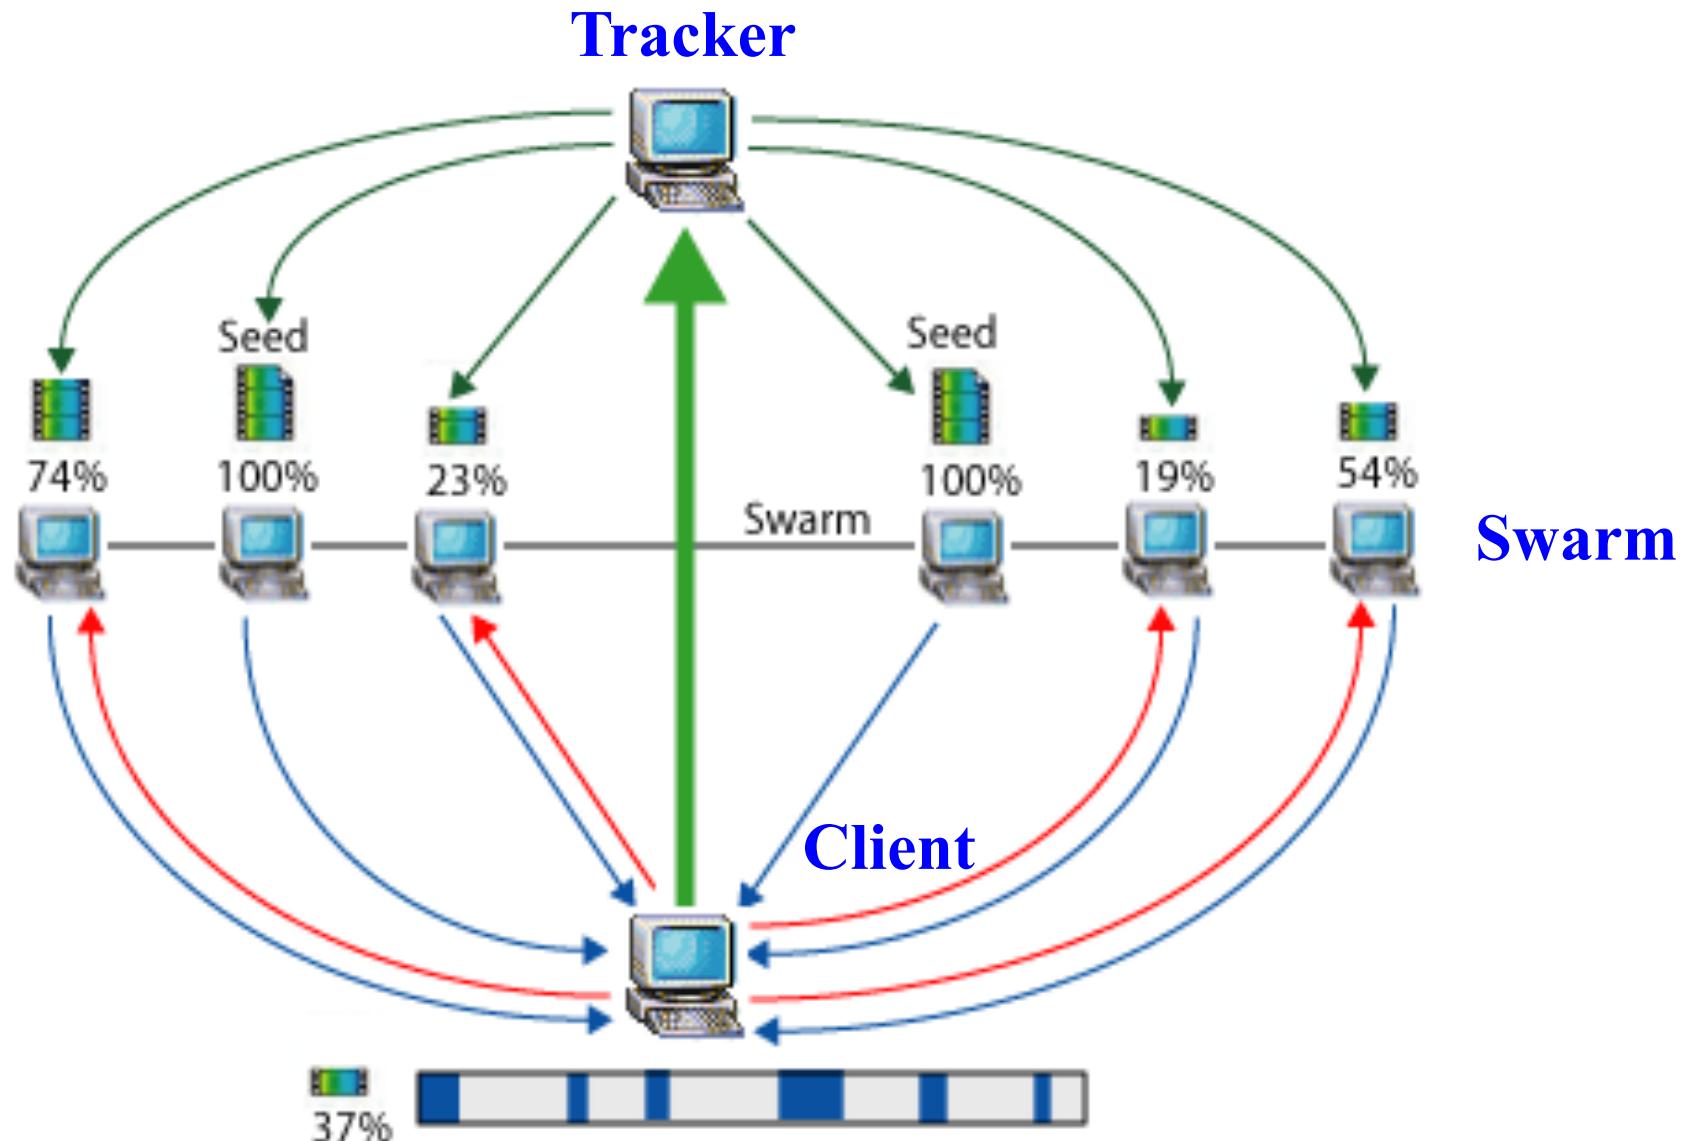
Agenda

- Internet update
- Short history of content delivery
- **Content delivery today**
- Network virtualization
- Challenges & opportunities

HTTP server selection



Server selection: P2P



World data centers



Google data centers



<http://royal.pingdom.com/2008/04/11/map-of-all-google-data-center-locations/>

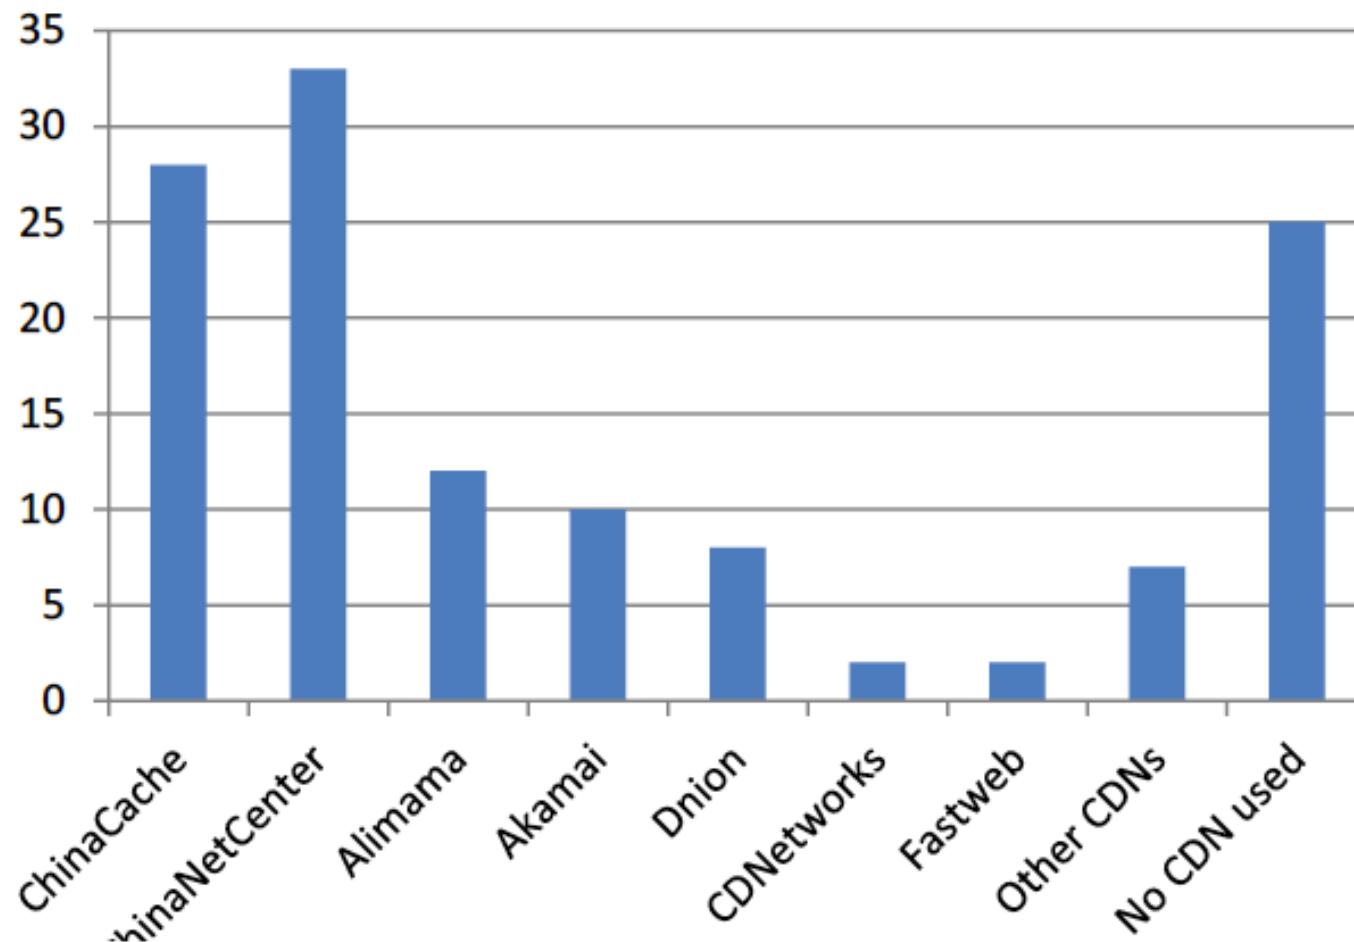
Where is content?

- California
- China already 2nd!
- USA: 9 among top 20
- Other developed countries: limited own content

Rank	Country	Potential	Normalized potential
1	USA (CA)	0.254	0.108
2	China	0.128	0.107
3	USA (TX)	0.190	0.061
4	Germany	0.183	0.058
5	Japan	0.163	0.051
6	France	0.146	0.034
7	Great Britain	0.157	0.030
8	Netherlands	0.144	0.029
9	USA (WA)	0.135	0.027
10	USA (unknown)	0.164	0.027
11	Russia	0.038	0.027
12	USA (NY)	0.130	0.026
13	Italy	0.122	0.018
14	USA (NJ)	0.125	0.016
15	Canada	0.028	0.015
16	USA (IL)	0.116	0.014
17	Australia	0.118	0.013
18	Spain	0.116	0.013
19	USA (UT)	0.111	0.012
20	USA (CO)	0.113	0.012

CDNs in China

- Select top 100 Chinese websites (Alexa)
- Which CDNs serve content from these websites?



Largest China Network by Coverage and Capacity

ChinaCache's content delivery network covers over 120 cities throughout China with over 350 nodes and 12,000 servers.

We're 15x larger in both capacity and reach compared to all other global content delivery network providers in China combined.





● China Telecom

● China unicom

● China mobile

● Education Network

● China Railcom

● Great Wall Broadband Network

The “P2P CDN”

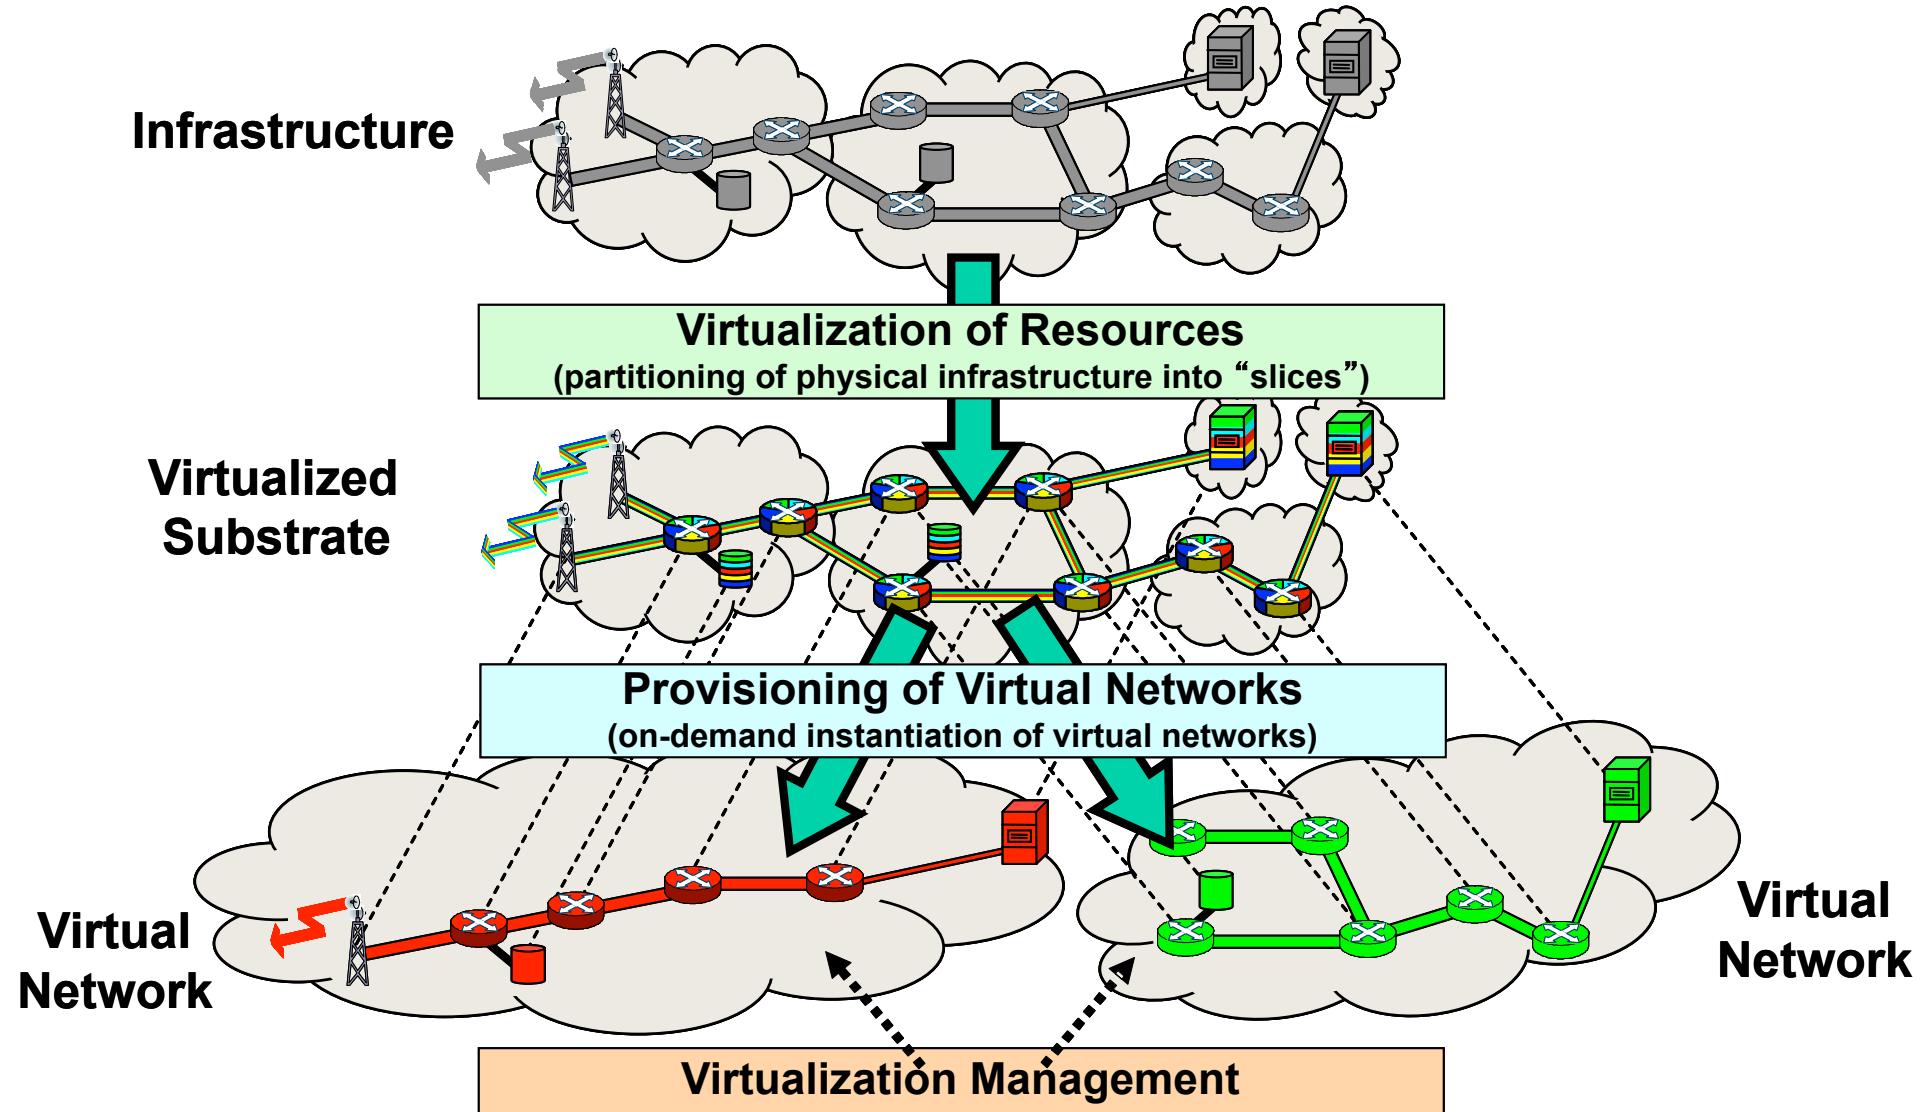
#	Country	U. Peers
1	US	1379462
2	ES	887480
3	GB	800308
4	CA	551820
5	IN	514246
6	AU	322009
7	BR	318294
8	IT	295339
9	PL	288780
10	PT	220124

TOP 10 ISPs (BT VIDEO USER)			
	AS#	Peers	AS Name-Internet Service Provider
1	3352	165469	TELEFONICA-DATA-ESPANA(TDE)
2	3662	129047	DNEO-OSP7-COMCAST CABLE
3	6461	127297	MFNX MFN-METROMEIDA FIBER
4	2119	113597	TELENOR-NEXTEL T.NET
5	19262	101390	VZGNI-TRANSIT-Verizon ISP
6	3301	97658	TELIANET-SWEDEN TELIANET
7	3462	96564	HINET-DATA CBG
8	4134	87392	CHINANET-BACKBONE 
9	6327	86964	SHAW-SHAW COMMUNICATION
10	174	74453	COGENT COGENT/PSI

Agenda

- Internet update
- Short history of content delivery
- Content delivery today
- **Network virtualization**
- Challenges & opportunities

The Virtualized Network



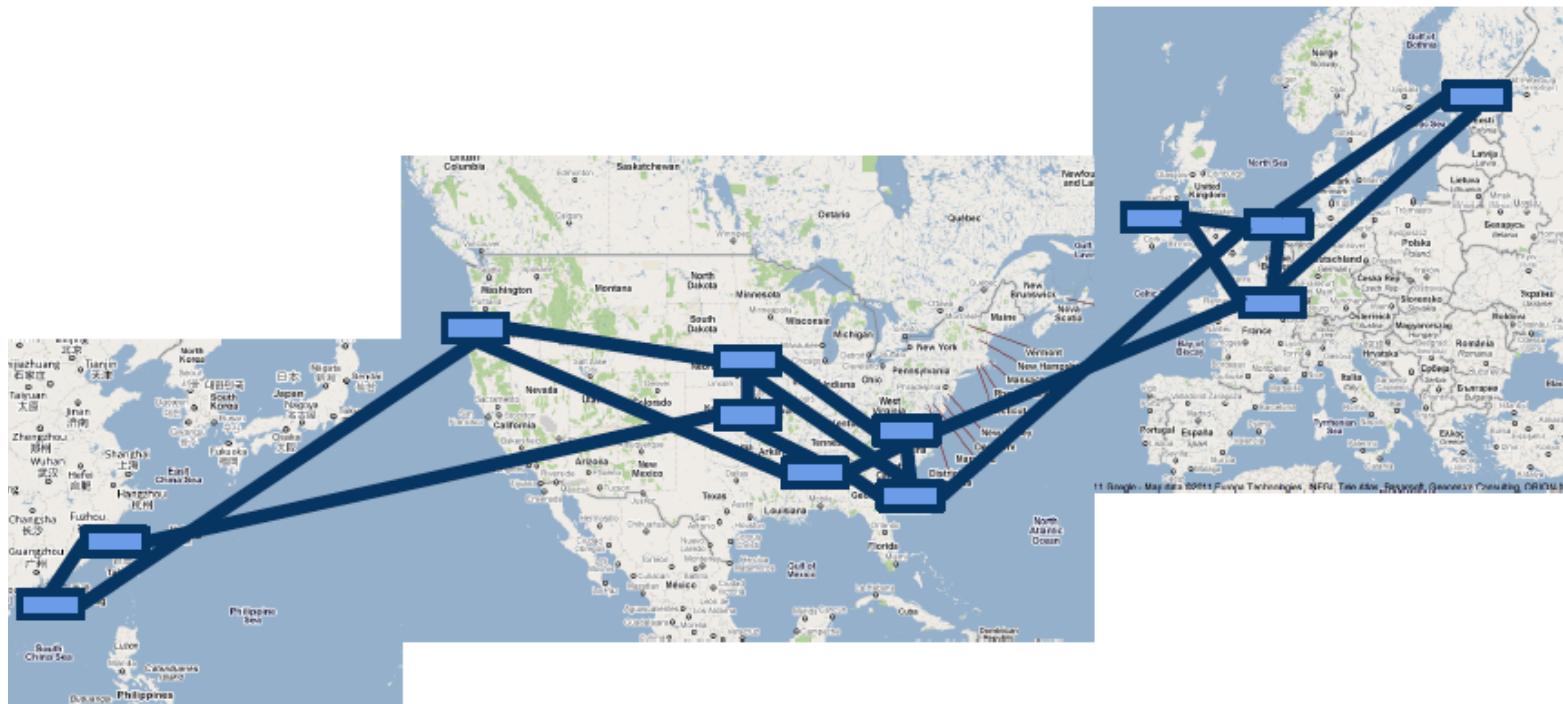
Virtual network = resource isolation/sharing

- Different architecture/protocol per virtual network
 - Does not have to be IP protocol, e.g., ICN
 - For QoS, security, different types of content/applications
- Expose network components to applications and services
- Dynamic: migration/expansion/contraction

Google and SDN

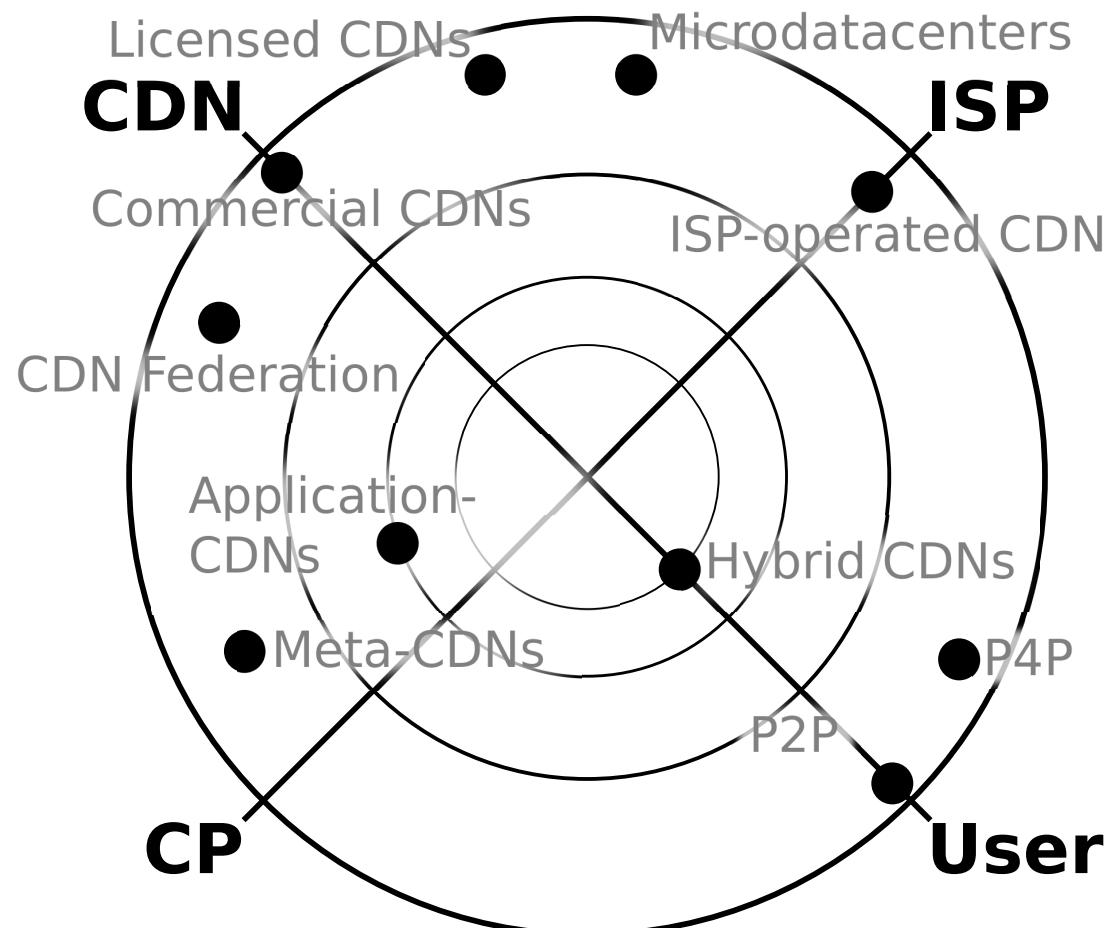
- Google is using OpenFlow
- Purpose: traffic engineering

Google's Software Defined WAN



Collaborative content delivery

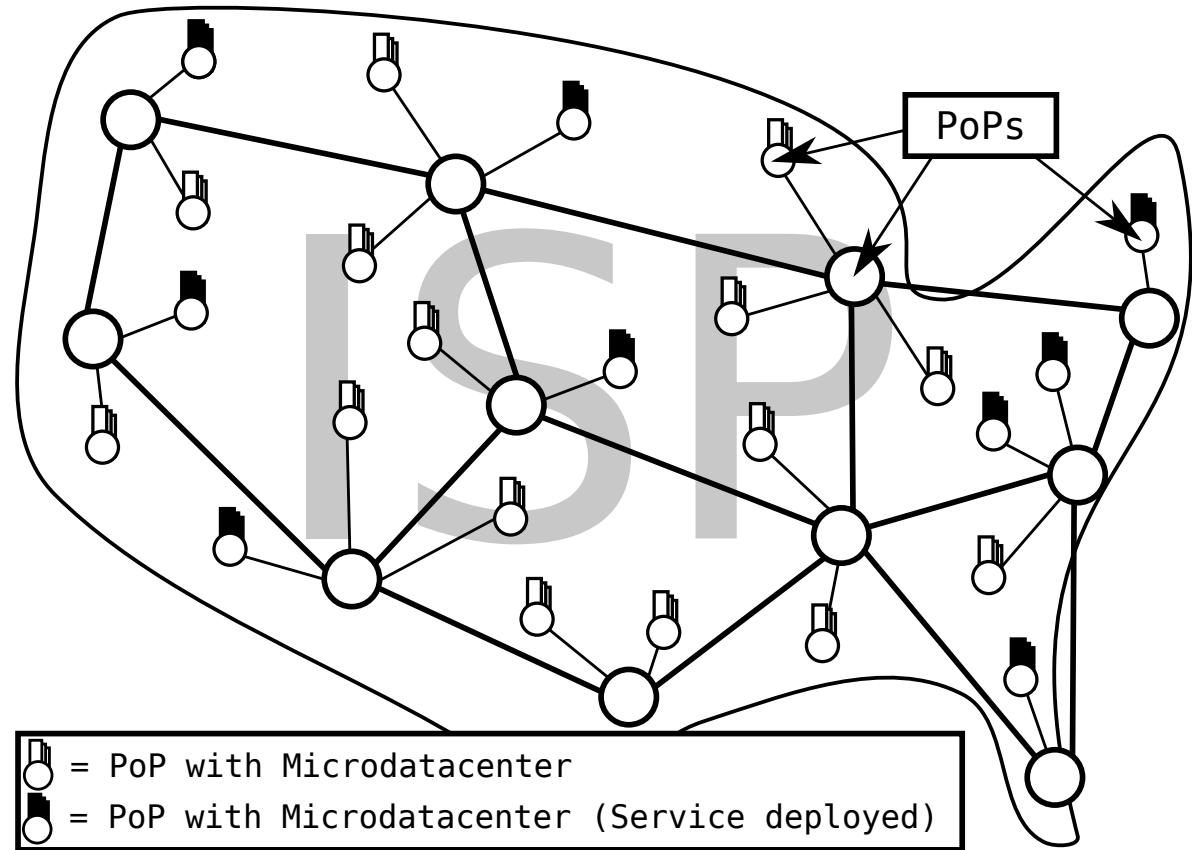
- Importance of stakeholders
- Spectrum in the solution space
- Enabling collaboration



B. Frank, I. Poese, G. Smaragdakis, S. Uhlig, and A. Feldmann. ***Pushing ISP-CDN Collaboration to the Limit.*** ACM SIGCOMM Computer Communication Review, 43(3):35-44, July 2013.

CDN 3.0

- Hybrid infrastructures: Akamai, PPTV
- Meta-CDNs, e.g., Conviva
- Virtual CDNs through ISP micro-datacenters
- ICN



B. Frank, I. Poese, G. Smaragdakis, V. Aggarwal, A. Feldmann, S. Uhlig, B. Maggs, F. Schneider. ***Collaboration Opportunities for Content Providers and Network Infrastructures***. SIGCOMM e-book chapter, to appear, 2013.

Agenda

- Internet update
- Short history of content delivery
- Content delivery today
- Network virtualization
- Challenges & opportunities

Challenges & opportunities

- Tracking the changing Internet ecosystem
- Scalability of massive content infrastructure
- Collaborative and agile content delivery
- Virtualization for increased diversity and agility

References

- C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. *Internet Interdomain Traffic*. Proc. of ACM SIGCOMM, 2010.
- I. Poese, B. Frank, B. Ager, G. Smaragdakis, and A. Feldmann. *Improving content delivery using provider-aided distance information*. Proc. of ACM SIGCOMM IMC 2010.
- B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig. *Comparing DNS resolvers in the wild*. Proc. of ACM SIGCOMM IMC 2010.
- B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig. *Web content cartography*. Proc. of ACM SIGCOMM IMC 2011.
- William B. Norton. *The Internet Peering Playbook : Connecting to the Core of the Internet*. DrPeering Press, 2012.
- B. Frank, I. Poese, G. Smaragdakis, S. Uhlig, and A. Feldmann. *Enabling Content-aware Traffic Engineering*. ACM CCR, 42(5):21-28, October 2012.
- B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. *Anatomy of a Large European IXP*. Proc. of ACM SIGCOMM, 2012.
- B. Frank, I. Poese, G. Smaragdakis, S. Uhlig, and A. Feldmann. *Pushing ISP-CDN Collaboration to the Limit*. ACM CCR, 43(3):35-44, July 2013.
- B. Frank, I. Poese, G. Smaragdakis, V. Aggarwal, A. Feldmann, S. Uhlig, B. Maggs, F. Schneider. *Collaboration Opportunities for Content Providers and Network Infrastructures*. SIGCOMM e-book, to appear, 2013.